16 allocated for parity information. As a consequence, the representations of data and parity information stored on Disk Drives 16 and illustrated in FIGS. 7A and 7B will appear as previously discussed herein above, except that Disk Drive(s) 16P will contain no Parity Blocks (PBs) 98 relating to the Data Blocks (DBs) 96 written while operating in the parity inhibited mode.

Thereafter, at the control of the user and generally when the rate at which new data must be written to Disk Drives 16 has returned to lower levels, the user may enter a command through User Input 102, identified in FIG. 8 as Parity Generate Assert (ParityGenA) 102B, directing the system to generate and store parity information for the Data Blocks (DBs) 96 which have been written to Disk Drives 16 without corresponding Parity Blocks (PBs) 98. In response, Demand Paging and Memory Management Facility (DPMM) 106 will first determine which Data Blocks (DBs) 96 have been written without corresponding Parity Blocks (PBs) 98, which may be accomplished in a number of ways, as will appreciated by those of ordinary skill in the relevant arts.

For example, Demand Paging and Memory Management Facility (DPMM) 106 may, when writing Data Blocks (DBs) 96 in the parity inhibited mode, mark a field of each Data Item Entries (DIEs) 112 in Address Translation Table 108 corresponding to a Data Block (DB) 96 that was written without a corresponding Parity Block (PB) 98, and may subsequently read the Data Item Entries (DIEs) 112 to identify those marked as not having corresponding parity information. In an alternate embodiment, Demand Paging and Memory Management Facility (DPMM) 106 may construct a parity inhibited data block table in Address Translation Table 108 containing identifications of each of the Data Blocks (DBs) 96 written while in the parity inhibited mode. In yet another embodiment, Demand Paging and Memory Management Facility (DPMM) 106 may search Data Item Entries (DIEs) 112 to locate any Data Address Translation Information (DATIs) 114 not having a corresponding Parity Address Translation Information (PATI) 116 and therefore corresponding to Data Blocks (DBs) 96 written in the parity inhibited mode.

Having identified the Data Blocks (DBs) 96 for which parity information is not stored in Disk Drive 16P, Demand Paging and Memory Management Facility (DPMM) 106 will read the identified Data Blocks (DBs) 96 from Disk Drives 16D in the manner described above with regard to read operations of mirrored and non-mirrored Data Blocks (DBs) 96. Demand Paging and Memory Management Facility (DPMM) 106 will also read all other Data Blocks (DBs) 96 that will be involved in generating the parity information to be stored in the Parity Blocks (PBs) corresponding to the identified Data Blocks (DBs) 96, referring to the above discussion of the combinations of Data Blocks (DBs) 96 represented by each Parity Block (PB) 98, and will pass the read Data Blocks (DBs) 96 to a Parity Generator 128.

Parity Generator 128, in turn, will generate the parity information to be stored in the Parity Blocks (PBs) 98 corresponding to the identified Data Blocks (DBs) 96 and Demand Paging and Memory Management Facility (DPMM) 106 will write the Parity Blocks (PBs) 98 to Disk Drive 16P in the manner discussed above.

It will be appreciated by those or ordinary skill in the arts that while Parity Generator 128 is illustrated in FIG. 8 as associated with Demand Paging and Memory Management Facility (DPMM) 106 in Host Computer 10, Parity Generator 126 may be implemented in a number of ways in the mass storage system of the present invention. For example, Parity Generator 126 may also be implemented in Disk Controller 24 or in Disk Platform 28.

It will be understood, therefore, that the user selectable parity mechanism of the present invention allows parity information to be generated and stored "off-line", that is, after the data has been received rather than while the data is being received, thereby significantly reducing the time required to receive and store data in circumstances wherein the time available to receive and store data is a critical factor or limitation in use of the mass storage system, and allowing data to be received and stored at a significantly higher rate.

It will also be understood that the data protection mechanisms of the present invention as described above allow a user to selectively tailor the level of protection in a mass storage system to meet the changing requirements of operation and use of the system by providing selectable levels of protection. For example, at the lowest level the system may be operated without mirroring and with parity generation and storing inhibited, thereby providing the maximum rate of data reception and storage but the lowest level of protection. At the next level, parity information may be generated and stored while the data is being received and stored, thus providing a higher level of protection but a lower data transfer rate. At the next level, the user may select intra-mirroring, wherein data that is selected and identified by the user to be additionally protected is written to logical units of the disk drives that have been selected and designated as mirrored logical units, thereby providing a still higher level of protection for the most critical or important data at the cost of some increase in storage space. Finally, the using the same mechanism as provided for intra-mirroring but designating all logical units for mirroring, the user may achieve full mirroring for all data stored in the system, but at the cost of significantly increased storage space.

While the invention has been particularly shown and described with reference to preferred embodiments of the apparatus and methods thereof, it will be also understood by those of ordinary skill in the art that various changes, variations and modifications in form, details and implementation may be made therein without departing from the spirit and scope of the invention as defined by the appended claims. For example, not only may a greater or lesser number of disk drives be used, but both the data and the parity information may be distributed across a plurality of disk drives, and at least some or all of the disk drives may be used to store both data and parity information. Therefore, it is the object of the appended claims to cover all such variation and modifications of the invention as come within the true spirit and scope of the invention.

What is claimed is:

1. In a mass storage mechanism for a system having mass storage devices for storing data and parity blocks respectively containing data and parity information wherein the system includes a host processor including memory and disk management facilities and a disk platform connected from the host processor and controlling a plurality of disk drive units comprising the mass storage devices, a protection mechanism providing user selectable levels of protection against data loss, comprising:

the plurality of disk drives for storing data in data blocks in storage segments of the disk drives and for storing parity blocks in storage segments of the disk drives, wherein

the storage segments of the disk drives are organized into at least two functionally separate logical units, and

each parity block contains parity information relating to the data stored in a plurality of corresponding data blocks wherein each one of the corresponding data blocks are located in a different one of the data disk drives.

a disk allocation mechanism for storing an identification of at least one of the logical units to be mirrored, and

a memory management mechanism for controlling operations of the disk platform for writing data blocks and parity blocks into the disk drives,

the memory management mechanism being responsive to the identification of a logical unit for mirroring all data blocks written into the designated logical unit by

writing a first copy of a data block assigned to a first storage address in a designated logical unit into the assigned storage address in the designated logical unit and

writing a second copy of the data block assigned to a storage address in a designated logical unit into a second storage address in the disk drives wherein

the second storage address is skewed with respect to the first storage address so that the second storage address is located in a disk drive separate from the disk drive containing the first storage address, and

writing at least one parity block containing parity information relating to the data block into the disk drives.

2. The protection mechanism of claim 1, wherein the data blocks are organized in the data disk drives in a striped configuration so that each sequentially addressed data block is located in a different sequential one of the disk drives.

3. In a mass storage mechanism for a system having mass storage devices for storing data blocks containing data and parity blocks containing parity information wherein the system includes a host processor including memory and disk management facilities and a disk platform connected from the host processor and controlling a plurality of disk drive units comprising the mass storage devices, a protection mechanism providing user selectable levels of protection against data loss, comprising:

the plurality of disk drives for storing data blocks in storage segments of the disk drives and storing parity blocks in storage segments of the disk drives, wherein

the storage segments of the disk drives are organized into at least two functionally separate logical units for storing data blocks wherein each logical unit includes storage segments on each one of the disk drives, and

each parity block contains parity information relating to the data stored in a plurality of corresponding data blocks wherein each one of the corresponding data blocks are located in a different one of the disk drives, and

a memory management mechanism for controlling operations of the disk platform for writing data blocks and parity blocks into the disk drives, wherein

the memory management mechanism is responsive to assertion of a parity inhibit command for writing only data blocks into the disk drives,

the memory management mechanism is responsive to de-assertion of the parity inhibit command for

determining which data blocks have been written into the disk drives without corresponding parity blocks written into a disk drive,

reading the data blocks written into the disk drives without corresponding parity blocks and all associated data blocks having a corresponding parity block corresponding to each of the data blocks written into the disk drives without corresponding parity blocks, and

generating and writing into a disk drive at least one parity block corresponding to each of the data blocks written into the disk drives without corresponding parity blocks.

4. The protection mechanism of claim 1 wherein the identification of at least one of the logical units to be mirrored is selectively entered by user command.

5. The protection mechanism of claim 3 wherein the parity inhibit command is asserted and de-asserted by user command.

6. In a mass storage mechanism for a system having mass storage devices for storing data and parity blocks respectively containing data and parity information wherein the system includes a host processor including memory and disk management facilities and a disk platform connected from the host processor and controlling a plurality of disk drive units comprising the mass storage devices, a method for providing user selectable levels of protection against data loss, comprising the steps of:

storing the data and parity blocks in storage segments of a plurality of disk drives, including

organizing the storage segments of the disk drives into at least two functionally separate logical units for storing data wherein each logical unit includes storage segments on each one of the disk drives, and

each parity block contains parity information relating to the data stored in a plurality of corresponding data blocks wherein each one of the corresponding data blocks are located in a different one of the disk drives,

storing an identification of at least one of the logical units to be mirrored in a disk allocation mechanism, and

by operation of a memory management mechanism for controlling operations of the disk platform for writing data blocks and parity blocks into the disk drives,

responsive to the identification of a logical unit, mirroring all data blocks written into the designated logical unit by

writing a first copy of a data block assigned to a first storage address in a designated logical unit into the assigned storage address in the designated logical unit and

writing a second copy of the data block assigned to a storage address in a designated logical unit into a second storage address in the disk drives wherein

the second storage address is skewed with respect to the first storage address so that the second storage address is located in a disk drive separate from the data disk drive containing the first storage address, and

writing at least one parity block containing parity information relating to the data block into a disk drive.

7. The method of claim 6, further comprising the step of organizing the data blocks in the disk drives in a striped configuration so that each sequentially addressed data block is located in a different sequential one of the disk drives.

8. In a mass storage mechanism for a system having mass storage devices for storing data blocks containing data and parity blocks containing parity information wherein the system includes a host processor including memory and disk management facilities and a disk platform connected from the host processor and controlling a plurality of disk drive units comprising the mass storage devices, a method for providing user selectable levels of protection against data loss, comprising the steps of:

storing the data blocks and the parity blocks in storage segments of a plurality of disk drives, including

organizing the storage segments of the disk drives into at least two functionally separate logical units for storing data wherein each logical unit includes storage segments on each one of the disk drives, and

each parity block contains parity information relating to the data stored in a plurality of corresponding data blocks wherein each one of the corresponding data blocks are located in a different one of the disk drives, and

by operation of a memory management mechanism for controlling operations of the disk platform for writing data blocks and parity blocks into the disk drives,

asserting a parity inhibit command, and, in response to the parity inhibit command, writing only data blocks into the data disk drives, and

de-asserting the parity inhibit command, and, in response to the de-asserting of the parity inhibit command,

determining which data blocks have been written into the disk drives without corresponding parity blocks written into a disk drive.

reading the data blocks written into the disk drives without corresponding parity blocks and all associated data blocks having parity information is a parity block corresponding to each of the data blocks written into the disk drives without corresponding parity blocks, and

generating and writing into a disk drive at least one parity block corresponding to the data blocks written into the disk drives without corresponding parity blocks.

9. The method of claim **6** wherein the identification of at least one of the logical units to be mirrored is selectively entered by user command.

10. The method of claim **9** wherein the parity inhibit command is asserted and de-asserted by user command.

* * * * *